

APPROXIMATED PER-FLOW RATE LIMITING

David R. Cheriton

5 ABSTRACT OF THE DISCLOSURE

A method and apparatus to limit the throughput rate of non-adapting aggressive flows on a packet-by-packet basis. Each packet of an input flow is mapped to an entry in a flow table for each output queue. The mapping is based on a subset of the packet's header data, giving an approximation of per-flow management. Each entry contains a credit value. On packet reception, the credit value is compared to zero; if there are no credits, the packet is dropped. Otherwise, the size of the packet is compared to the credit value. If sufficient credits exist (i.e., size is less than or equal to credits), the credit value is decremented by the size of the packet in cells and the processing proceeds according to conventional methods, including but not limited to those disclosed in the co-pending DBL Application, incorporated herewith by reference in its entirety. If, however, the size of the packet exceeds the available credits, the credit value is set to zero and the packet is dropped. A periodic task adds credits to each flow table entry up to a predetermined maximum. The processing rate of each approximated flow is thus maintained to the rate determined by the number of credits present at each enqueueing decision, up to the allowed maximum. The scheme operates independently of packet flow type, providing packet-specific means for rapidly discriminating well-behaved flows that adapt to congestion situations signaled by packet drop from aggressive, non-adapting flows and managing throughput bandwidth accordingly. Bandwidth is shared fairly among well-behaved flows, large and small, and time-critical (low latency) flows, thereby protecting all from non-adapting aggressive flows.